

Toward a crew-system concept for real-time fault management in next-generation aerospace vehicles

Robert S. McCann, Ph.D.
NASA Ames Research Center
Moffett Field, CA

&

Jeffrey McCandless, Ph.D.
San Jose State University Foundation
Moffett Field, CA

Abstract

Real-time health management on today's shuttle missions is both difficult and labor-intensive, posing a significant risk to crew safety and mission success. Next-generation space transportation vehicles are expected to incorporate several advanced information processing and information display technologies, transforming fault management into a cooperative venture between crews and intelligent systems. Optimizing the potential of this crew-system partnership poses a considerable design challenge. The new technologies vastly increase the space of design options for knowledge engineering architectures, human-machine function allocation, and human-computer interfaces. We propose an integrated program of systems-level simulation, real-time human-in-the-loop scenario simulation, and human performance modeling to "prune" the design space and optimize the crew-systems concept.

Introduction

Designed over a quarter of a century ago, NASA's space transportation system is enormously expensive and operationally risky. Each shuttle mission costs approximately half a billion dollars, and probabilistic risk assessments have estimated the probability of catastrophic failure at somewhere between 1 in 250 to 1 in 450 missions. The combination of high cost and high risk constitute major barriers to broadening the scope of human activity in space. Recently, NASA initiated several high profile programs to develop reusable launch vehicles (RLVs) that are dramatically cheaper and more reliable than the shuttles. Nearer term, the Space Launch Initiative (SLI) is investing in component technologies to reduce the probability of a catastrophic vehicle failure in a second generation RLV to 10^{-4} per flight, at a launch cost of \$1000 per pound. Longer term, the third generation RLV program is targeting a vehicle with a catastrophic failure rate of only 10^{-6} , with per-mission launch costs of one hundred dollars per pound. Achieving these ambitious levels of reliability and efficiency is a major challenge. Human-bearing spacecraft consist of a large number of highly complex engineering systems. These include multiple propulsion systems, an electrical power generation and distribution system, a life support and environmental control system, a data processing system, and a communications system. Regardless of the thoroughness of ground-based maintenance procedures, the complex and dynamic nature of these systems raises the distinct possibility of a system malfunction during flight. Together, the extreme environments in which spacecraft operate, the toxic and physically volatile nature of propellants and other systems components,

and the high level of system interdependence create an environment in which system malfunctions can escalate rapidly into mission-threatening or life-threatening situations.

Shuttle crews deal with the danger of systems malfunctions by monitoring the health of the onboard systems via system summary displays and, in consultation with ground personnel, working systems malfunctions when they occur. Astronaut training places a strong emphasis on learning the architecture of the onboard shuttle systems, including their interconnections, interdependencies, and failure modes. This in-depth knowledge gives the crew the capability to troubleshoot malfunctions in real time. Even with the training, however, real-time fault management (FM) is a difficult and time-consuming activity. To illustrate these difficulties, we will describe the sequence of events that would unfold if a leak occurred in the external tank (ET) on ascent (powered flight).

To understand the effects of this leak, some details on the structure of the ET are necessary. The ET consists primarily of a forward tank containing liquid oxygen (LO₂) and an aft tank containing liquid hydrogen (LH₂). These propellants are fed to the three space shuttle main engines (SSME's) for combustion. As the supply of LH₂ is consumed, a cavity is created in the aft tank that must be pressurized to keep the flow within nominal levels. Cavity (ullage) pressure is maintained by a return flow of gaseous H₂ (GH₂) from the SSMEs. Each engine's GH₂ feedline has its own flow control valve (FCV) and its own dedicated ullage pressure sensor in the LH₂ tank. Normally, ullage pressure is controlled automatically: when the sensed pressure falls below a designated lower value, the corresponding FCV valve is opened; when the value returns to normal range, the FCV is closed. Note that this logic works on the three flow control valves (one for each SSME feed) independently, so the actual valve configuration is dynamic over time.

The primary symptom of a leak in the aft tank is abnormally low ullage pressure. Each sensor's current pressure reading is displayed to the crew on a systems summary display called the BFS SYS SUM 1 display. If any of the three sensors returns a value below the automatic set point (consistent with a leak), a "down" arrow appears beside the low reading on the display, and a class 3 alarm is issued by the shuttle's caution and warning system. The alert consists in part of a short tone and the appearance of a fault message on both the BFS GNC SYS Sum 1 display and on a dedicated fault message summary display. Therefore, the crewmembers' first task, after being alerted to the presence of a malfunction by the Class 3 alarm, is to read the fault message. In this particular case the message is a rather cryptic "MPS LH₂ ULL", indicating a low or high ullage pressure condition in the aft tank.

Low ullage pressure is a potentially dangerous situation; it can lead to cavitation in the fuel turbopumps and uncontained SSME damage. The goal for the crew, then, is to execute reconfiguration procedures as quickly as possible to restore ullage pressure to its nominal range. On ascent, malfunction procedures are located in an ascent/entry systems procedures (AESP) data file, one of several paper flight data files carried onboard the orbiter. The "low ullage pressure" section in the AESP contains the following instructions:

" if 2(3) Ps < 21.6 or > 34.5;
 MPS LH₂ ULL PRESS - OP
 When all Ps > 34.5:
 MPS LH₂ ULL PRESS - AUTO"

The first instruction is significant. It tells the crewmember that before effecting any reconfiguration, he or she must first check the BFS SYS SUMM 1 display to confirm that at least

two of the three sensed values are showing off-nominal readings. This cross-check is necessary because the caution and warning system cannot distinguish between real faults and off-nominal readings due to failed sensors. If only one of the three readings is low, the problem is likely a failed sensor rather than a real malfunction.

Once the malfunction is confirmed, the crewmember moves on to (in this case) reconfigure the affected system to optimize existing capabilities. As we see from the second line of the AESP entry, the appropriate procedure is to toggle the LH2 ULL PRESS switch (located on panel R2) from "Auto" (the nominal setting) to "OP" (for open). This switch throw opens the flow control valves from all three main engines simultaneously, maximizing the flow of GH2 into the aft tank. Then, once the switch has been toggled, the crewmember is told to monitor the ullage pressure readings on the BFS SYSS SUM 1 display to confirm that they are returning to the nominal range. Once this is confirmed, the final instruction is to toggle the LH2 ULL PRESS switch back to "AUTO".

The ULL PRESS example illustrates several generic difficulties with FM on the shuttle. Malfunction confirmation (the check that at least two of three sensors are showing off-nominal values, in this case) is a standard requirement that can be very demanding. Malfunction effects often propagate both within and between systems through interconnections and interdependencies. A second route for between-systems propagation is simple physical proximity between failed system components. Thus, malfunctions are frequently accompanied by a cascade of alarms and fault messages. When this occurs, the crew and/or MCC personnel must make a "root-cause" fault determination, which may include resetting parameter values to determine if the malfunction re-occurs in a timely manner, checking for related systems malfunctions, and performing inventories of failed equipment.

Inefficiencies in the fault management process, particularly those that lengthen the time course of the process, pose the biggest danger during ascent and entry. Faults in the main propulsion system, in particular, have short time frames before they constitute a hazard, putting a premium on executing the appropriate procedures as quickly as possible. Furthermore, in laboratory studies, human operators of complex physical systems exhibit a stubborn performance limitation known as cognitive lockup (Moray & Rotenberg, 1989; Moray, Inagaki, & Itoh, 2000). Briefly, when a system malfunction occurs in a high-workload environment, operators tend to focus exclusively on the fault management activities tied directly to that malfunction; they show very limited ability to time-share those activities with any concurrent information processing requirements. On ascent, shuttle crews have a well-defined set of display monitoring activities to maintain situation awareness of the vehicle's navigation state and SSME functioning. Cognitive lockup suggests that these ongoing activities will be omitted (or curtailed) during active fault management (Moray, et al.). Thus, we would expect the crews' situation awareness of overall vehicle state to degrade during a fault management exercise, and the longer the process takes, the more severe the degradation should become. Cognitive lockup is also likely to degrade the crew's ability to manage multiple fault scenarios when fault management activities overlap. To the extent that the current inefficiencies in the fault management process extend the time course of the process, they increase the probability of overlap.

Fault management on next-generation vehicles

Stringent safety and reliability requirements for next-generation spacecraft mandate a dramatic improvement in the fault management process compared to the shuttle. Fortunately, designers of the new vehicles have almost thirty years of advances in the fields of computer

science and information technology to draw on. At the sensor level, integrated vehicle health management (IVHM) programs have developed (or are developing) a new generation of sensors that are more reliable, and capable of sensing more forms of data (e.g., plume composition, vibration) than the sensors on the shuttle. These "smart" sensors promise to provide a much more complete and reliable picture of real-time systems functioning than on the shuttles. Second, artificial intelligence researchers have developed sophisticated inference engines that directly support real-time vehicle health management. For example, model-based reasoning agents (Patterson-Hine, Hindson, Sanderfer, Deb, & Domagala, 2001) take data feeds from sensors, compare the data to a structured model of subsystem components, their functions, and interconnections, and make inferences about system state. Nominal data patterns are consistent with a nominal configuration of all subsystem components. If a data pattern goes off-nominal, the reasoning agent infers the most likely state change that would produce the off-nominal pattern, such as a valve failed open or failed closed. From there, intelligent inferences can be made as to the appropriate reconfiguration procedure.

NASA's next generation vehicle development programs are targeting technologies to further increase the power and reliability of the inference engines. One of the most promising areas is in knowledge fusion, whereby model-based reasoning is combined with other forms of artificial reasoning (rule-based, neural network pattern recognizers, etc.) to enhance the accuracy and reliability of nominal and off-nominal state determinations. Moving beyond the systems level of analysis, another focus is on developing intelligent crew assistants. Some of the capabilities discussed for these flexible agents are performing root-cause determinations for faults that have intra-system impacts, analyzing impacts of malfunctions on vehicle functioning and mission goals, and performing real-time mission replanning.

Incorporating these advanced information technologies creates two critical design issues. One is to define an optimal knowledge engineering architecture. The availability of single-chip control and communication processors equipped with real-time operating systems has made it possible to embed powerful data processing algorithms all the way down to the sensor level. Higher level controllers will reside on machines with vast computational power. The challenge is how to best distribute information-processing and decision-making capabilities between the various levels of the hierarchy.

One architectural candidate is a three-tiered pyramid structure, with controllers at the subsystem (sensor) level, the system level, and the vehicle level. Each level would integrate information from relevant lower levels and exploit an increasingly broad knowledge base and decision-making capability. For example, returning to our low ullage pressure case, the three sensors in the LH2 tank might feed their output to a local controller that checks and compares the data values. One quite useful function for the controller would be to implement a command voting scheme similar to the scheme employed by the SSME's digital computer units. If two of the three sensor outputs are identical, and the third is anomalous, the third sensor is declared failed, and its data discarded. In this way, off-nominal annunciations based on a single sensor value would be suppressed.

The second or systems-level controller might process and integrate information from the subsystem controllers. Model-based reasoning agents are obvious candidates for this level, as they are designed to integrate sensor data from component subsystems to make system-wide state determinations. And finally, a single vehicle level controller would essentially function as the intelligent crew assistant mentioned earlier.

The second major design issue concerns the role of the human in this brave new world of distributed intelligent systems. At first glance, the issue might seem moot. Several model-based reasoning agents have demonstrated an ability to perform fault detection, isolation, and recovery (FDIR) activities autonomously. Thus, at least in principle, the labor-intensive fault management activities required of the crew and mission control experts on the shuttle today could be performed by machine intelligence. However, automated systems are never going to achieve 100% reliability. They can fail outright if their hardware platform is damaged, and as memory chips shrink to ever-smaller sizes, space-based software become increasingly vulnerable to single event upsets (individual bit flips) from contact with charged particles. Add the fact that automated reasoning agents are unlikely to achieve 100% accurate state assessments in all cases, and it is clear that humans must retain a FM role.

Subject matter experts who have studied the impact of automation in transport aircraft are strongly opposed to the concept of full automation (Billings, 1997). As applied to fault management, giving full responsibility to the machine would change the crew's role from an active participant in the process to an (at best) passive monitor of the process. This change in crew role carries several human factors risks. If the crew is not actively involved in decision-making, they can become overreliant on the automation, and fail to perform the ongoing information sampling activities needed to maintain a high level of situation awareness concerning system state and system functioning, knowledge that they would need in case the automation failed. A second potential problem with automating all FM activities is mode confusion, a lack of real-time understanding of what the automation is doing and/or why it is doing it. For example, suppose an automated fault management system diagnosed a malfunction and started executing the relevant procedures, but did not keep the crew well informed of what it was doing. Serious problems could arise if the crew made a malfunction diagnosis that disagreed with the automated diagnosis, and attempted a system reconfiguration that cancelled or reversed the automated action.

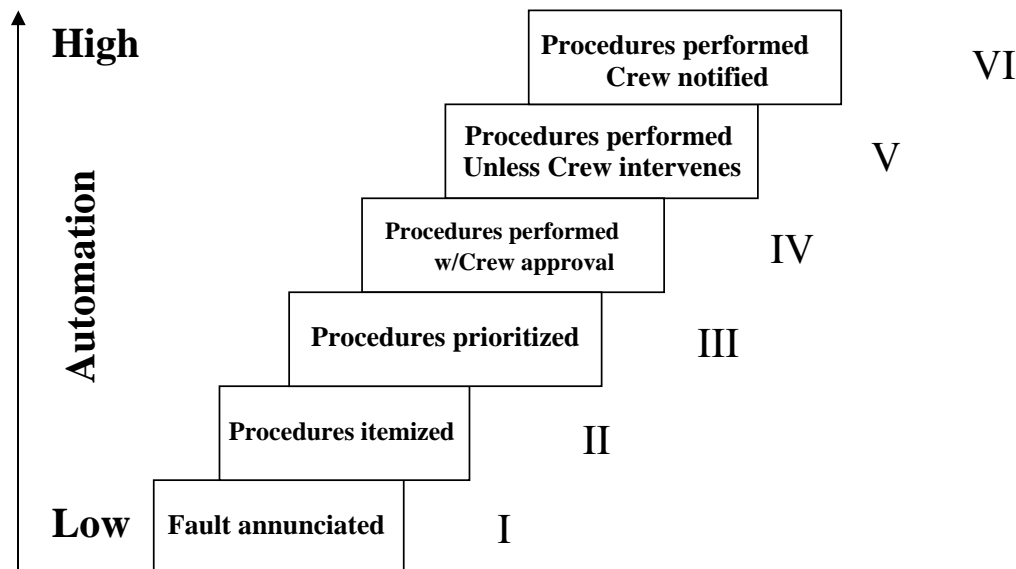


Figure 1. Possible levels of automation for fault management

Billings (1997) argued that the best way to retire these risks is to adopt a human-centered approach to human-machine function allocation. Ideally, automation and human would form a cooperative team with a functional allocation that keeps the human firmly in the loop, while still maximizing the capabilities of the intelligent systems. A useful framework within which to explore this concept is a FM version of the standard Sheridan-Verplaank (S-V) levels of automation (LOA) scale. Shown in Figure 1, the scale defines several forms of human-machine function allocation. At the bottom of the scale, corresponding to the highest level of human involvement, the machine announces a fault and leaves all other FM responsibilities to the human. This is roughly the situation on the shuttles today. Intermediate levels assign most of the fault detection and fault diagnostic responsibilities to the machine, together with the ability to provide the appropriate list of procedures to the human in electronic form. The human still performs the reconfiguration procedures. At the higher levels, automation assumes all responsibilities, including reconfiguration. The human still has an important function, however. In S-V level IV, the machine does not perform reconfiguration procedures until the human gives permission; in level V, the machine carries out the procedures unless the human vetoes the actions within a certain period of time. Finally, the highest LOA gives the machine control over all FM activities, informing the human only after procedures are complete.

Which LOA is the optimal choice to satisfy human-centered and system-centered requirements? The short answer is that no single level is appropriate for all conditions. Figure 2 plots a hypothetical LOA curve as a function of two critical variables: crew workload (which is highly correlated with flight phase), and time criticality of the malfunction. The nonlinear shape of the LOA curve points to the relative importance of time criticality to the LOA determination. The hashed rectangle overlaying the top right portion of the function is a particularly important section. It represents the set of highly time-critical malfunctions that evolve too rapidly for humans to play a useful role. These faults obviously mandate S-V level VI, the highest possible. Below this boundary, the gradual reduction in the slope of the curve reflects the increasing influence of crew workload as more time is available. We noted earlier that if a




Figure 2. Adaptive automation as a function of time criticality and workload

crewmember is working a fault management problem, he or she might postpone work on a second malfunction until the first is completed. Thus, the same fault might be assigned a higher LOA if it occurs second in a two-failure deep scenario, compared to when it occurs first. An alternative design strategy might be to assign a higher LOA to first failures, thereby saving the human to devote his or her full attention to a second fault if, in general, later faults tend to be more time-critical than earlier faults.




Figure 3. Freon loop under off- nominal conditions prior to automated action.

Most faults fall into an intermediate range of time criticality and operator workload, where the most appropriate LOA is probably one of the intermediate cases, where the automation diagnoses and verifies the malfunction and performs the appropriate procedures, but the human has "right-of-refusal" control over the computer's actions. There is evidence (Boy, 1988) that crew-systems might perform particularly well if the automation provided an initial diagnosis, and carried out the associated procedures, while the human engaged in "backward chaining" to verify the computer's diagnosis against the actual fault symptoms.

To flesh out the user interface to support this form of human/machine functioning, a new generation of systems summary displays will have to be designed. Similar to those on the MD-11 today, these displays will likely integrate the traditional concept of a systems summary display with the procedural knowledge contained in the flight data files. The goal would be a display that explicitly depicts the automated actions in a form that is inspectable, predictable and isomorphic with the operator's mental model of the affected system (Weiner, 1989).

One display concept that holds considerable promise for this area is the animated mimic display (Bennett & Malek, 2000), a graphical representation of a complex system that provides a visual depiction of the dynamics of system functioning. For example, suppose an intelligent fault management agent was developed for the shuttle, and accurately diagnosed an "Evap out T




Figure 4. Freon loop with automation reconfiguring the flow proportion valve setting.

low" problem in the environmental control and life support system. This problem occurs in the freon loop when freon temperature drops below 32 degrees F after being cooled by the flash evaporator subsystem. "Evap out T low" commands a sequence of procedures that must be worked in a matter of seconds to prevent the abnormally cold freon from freezing the water in the freon/water heat exchanger, located downstream of the flash evaporator. A display concept for the initial procedure is illustrated in Figures 3 and 4, which provide a "zoomed-in" view of the portion of the freon loop containing the "evap out T" sensor area and salient subsystems downstream. One of these is a flow proportion valve that controls the amount of freon flowing to the Water/Freon Interchanger (ICH) and the Payload Heat Exchanger (labeled PL). Since the immediate danger is that water will freeze in the ICH, the first procedure is to switch the flow proportion valve from ICH to PL, which diverts most of the flow away from the interchanger. In Figure 3, the Evap out T low condition has been identified (cf. the message in lower right), but no action has yet been taken by the automation.

In Figure 4, the freon/water loop interchanger and the flow proportion valve are highlighted in grey. The effectors appear on the display as virtual switches, and "flip" in real time to the PL position. As the actual flow proportion shifts in response to the reconfiguration, the size of the arrow illustrating real time flow volume from the flow proportion valve to the payload HX increases, while the arrow illustrating flow to the ICH decreases. Highlighting of this sort would track the actions of the automation precisely, giving the crew salient perceptual feedback as to what procedure was being worked, what actions were being taken, and how these actions were impacting real-time system functioning. A further attractive design feature would be to color code or otherwise identify the subsystem that was next scheduled for reconfiguration.

Pruning the design space with modeling and simulation

The various aspects of fault management discussed so far highlight the design challenge that accompanies the integration of advanced intelligent systems into the cockpit. Not only is the knowledge engineering design space itself quite extensive, but as we have seen, intelligent systems generate several options for human/machine functional allocation (Figure 1), each of which may require a customized user interface. The design space problem worsens when we consider that the spaces are not independent. When a fault occurs, it will be up to the intelligent system to evaluate crew workload, fault criticality and other relevant aspects of the real-time environment, and select an appropriate LOA. This is a critical decision: for example, next-generation vehicles are likely to be equipped with crew-escape modules designed to eject in the event of a catastrophic vehicle malfunction. The go/no go ejection decision is going to take place under extreme time constraints, and it will be up to the intelligent system to determine whether to eject automatically, or to allow the crew some role in the decision making process. Lives will depend on this decision. Less drastically, from the crew's perspective, intelligent systems will have to function as integrated signal detection devices, making decisions as to when an off-nominal condition is occurring, whether to initiate a caution and warning annunciation, and what level of criticality to assign to the alert.

These interconnections between the crew, the environment, and intelligent systems inflates the design space in what amounts to a combinatorial explosion. We must therefore find some way to prune this design space and identify the optimal design options. We propose an integrated program of modeling and simulation involving virtual vehicle systems models, human-in-the-loop simulations, and cognitive modeling. A brief description of this concept follows.

Virtual systems modeling. The nascent field of virtual engineering could be of great assistance to the crew-systems design process. The basic idea is to use high-speed numerical computing techniques, such as those used today to model the impact of high-speed collisions on automobile structures, to develop systems models capable of simulating dynamic aspects of systems functioning (e.g., flow rates, pressures, vibrations, thermodynamics) in real time. By incorporating virtual sensors into these models, real-time data feeds could be generated and provided to system-state inference engines. Assuming the models can incorporate the effects of systems malfunctions, Monte-Carlo style simulations could be run to train the inference engines to recognize (classify) a wide variety of system states. Measures of model performance could then be used to answer important IVHM questions, such as how many sensors to build into a system, and where they should be located.

Human-in-the-loop simulation. Real-time systems models could also form the backbone of a high-fidelity testing and evaluation simulation facility. Malfunction scenarios could be implemented that provide integrated tests of candidate knowledge engineering architectures, user interfaces, and human-machine function allocation schemes. To illustrate the need for this approach, earlier we noted some compelling reasons why intelligent systems should select the LOA on a malfunction-by-malfunction basis (adjustable automation). However, an AI system capable of exercising this level of control is virtually guaranteed to produce occasional emergent behaviors not anticipated by the designers, and that will come as a surprise to the humans. Humans do not respond well to unpredictable automation, and by measuring their response to these unexpected behaviors, we can determine whether the behavior is unacceptable, requiring modifications to the knowledge architecture.

A further justification for human-in-the-loop simulation is to determine the boundary conditions for the highest LOA (VI) determinations. We have already noted the need for a precise judgment of the temporal boundary that distinguishes malfunctions that are too time

critical to permit human involvement in the FM process from malfunctions that are not. This assessment cannot be made in the abstract. It will depend on the specifics of the user interface design, as the interface influences how quickly the human can recognize the fault and consider its impact. The assessment will also depend on the prognostic capabilities of the intelligent systems, which will determine how much lead-time exists between fault detection and catastrophic failure. Integrated tests of human/machine functioning can provide an empirical resolution of this issue.

Cognitive modeling. A third form of modeling that could contribute to pruning the space of crew-system design options is human cognitive modeling. Cognitive models take a particular task environment (such as fault management) and carefully decompose the task into its constituent elements. Based on their knowledge of human perceptual, cognitive, and motor systems functioning, the models make assessments of the time required to complete each element. Applied to the crew-systems realm, cognitive models have the potential to provide accurate assessments of the time course of human task management activities, including information acquisition, reasoning, and system reconfiguration operations. Human-in-the-loop simulations are expensive and time consuming, and only a small subset of the parameter space can be examined in such studies. Cognitive models may allow us to make quantitative assessments of human/machine performance over a much wider range of design options than human-in-the-loop simulation alone.

References

- Bennett, K. B., & Malek, D. A. (2000). Evaluation of alternative waveforms for animated mimic displays. *Human Factors*, 42, 432-450.
- Billings, C. E. (1997). *Aviation automation: The search for a human-centered approach*. Hillsdale, NJ: Erlbaum
- Boy, G. A. (1988). Operator assistant systems. In Hillnagel, E., Mancini, G., & Woods, D.D. (Eds). *Cognitive engineering in complex dynamic worlds*. New York: Academic Press (pp. 85-98).
- Moray, N., Inagaki, T., & Itoh, M. (2000). Adaptive automation, trust, and self-confidence in fault management of time-critical tasks. *Journal of Experimental Psychology: Applied*, 6, 44-58.
- Moray, N., & Rotenberg, I. (1989). Fault management in process control: Eye movements and action. *Ergonomics*, 32, 1319-1342.
- Patterson-Hine, A., Hindson, W., Sanderfer, D., Deb. S., & Domagala, C. (2001). *A model-based health monitoring and diagnostic system for the UH-60 Helicopter*. Presented at the American Helicopter Society 57th annual forum. Washington, DC.
- Weiner, E. L. (1989). *Human factors of advanced technology ('glass cockpit') transport aircraft*. NASA Tech. Report #117528, NASA Ames Research Center, Moffett Field, CA.